

1

SWITCH PROXY CONTROLLER FOR SWITCH VIRTUALIZATION

CROSS-REFERENCE TO RELATED APPLICATION

This application is a continuation of International Application No. PCT/US2018/064000, filed Dec. 5, 2018, which claims the benefit of U.S. Provisional Application No. 62/594,986, filed Dec. 5, 2017, and are both incorporated by reference in their entireties.

BACKGROUND

Switches may be used to route network traffic in a data center. For example, the switches may route packets from a particular Internet Protocol (IP) address to another IP address, block packets from another particular Internet Protocol (IP), and bandwidth throttle packets from yet another particular Internet Protocol (IP) address to another IP address. These switches are generally programmed by controllers to route network traffic.

SUMMARY

This document describes techniques, methods, systems, and other mechanisms for switch virtualization by a switch proxy controller. Having multiple controllers program a switch may be desirable in datacenter operation and management. However, this presents a problem because switch controller protocols and Application Programming Interfaces (APIs) typically expect only a single master controller to program a switch. Controllers are built with the assumption that they have exclusive access to program a switch and are not designed to co-exist with other controllers simultaneously programming the switch.

A system that includes a switch proxy controller may enable multiple controllers to program a single switch. The multiple controllers may each correspond to different switch fabrics. For example, a switch may be simultaneously controlled by both a first switch fabric and a second switch fabric. A switch fabric may include multiple instances of a controller where a single instance in the switch fabric is designated as a master and is the only instance in the switch fabric that the switch permits to program the switch. The first switch fabric may refer to a first switch fabric controller and a second switch fabric may refer to a second switch fabric controller.

The switch proxy controller may cause the switch to believe that it is being controlled by a single controller designated as master in a single switch fabric, and cause each of the multiple controllers designated as a master in the different switch fabrics to believe that they have exclusive access to program the switch. In doing so, the switch proxy controller may receive requests from the multiple controllers, provide requests to the switch, and provide responses to the requests received from the switch to the respective controller that provided the request. Accordingly, the switch proxy controller may enable the switch to behave as if only a single controller were providing programming instructions even when multiple controllers are controlling the switch.

Additionally, the switch proxy controller may hide information provided by the switch so that each switch fabric is unaware of information corresponding to the other switch fabric. For example, the switch proxy controller may receive a request for active routing rules on a switch from a particular controller, provide the request to the switch,

2

receive a list of all routing rules from the switch, identify the routing rules in the list that were programmed by the particular controller, and then provide just those routing rules that were identified to the particular controller without providing the routing rules that were programmed by other controllers.

Accordingly, the switch proxy controller may enable each of the controllers to be unaware that other controllers are also programming the switch so that each controller functions as if it had sole control of the switch. With the functionality described above, the switch proxy controller may ensure protocol version compatibility between switch fabrics and a switch, filter programming requests from the switch fabrics to prevent conflicts and resource exhaustion on the switch, and relay events from the switch to multiple switch fabric.

One innovative aspect of the subject matter described in this specification is embodied in a method that includes the actions of receiving, by a switch proxy controller, a first request from a first switch fabric, where the first request indicates a first identifier that identifies the first request from other requests from the first switch fabric, generating, by the switch proxy controller, a second request that indicates a second identifier that identifies the second request from other requests sent from the switch proxy controller to a switch, providing, by the switch proxy controller, the second request to the switch, receiving, by the switch proxy controller in response to the second request, a first reply that indicates the second identifier indicated in the second request, generating, by the switch proxy controller and based on the second identifier indicated in the first reply, a second reply that indicates the first identifier, selecting, by the switch proxy controller, the first switch fabric to receive the second reply based on the second identifier, and providing, by the switch proxy controller, the second reply to the first switch fabric.

Other embodiments of this aspect include corresponding computer systems, apparatus, and computer programs recorded on one or more computer storage devices, each configured to perform the actions of the methods. A system of one or more computers can be configured to perform particular operations or actions by virtue of having software, firmware, hardware, or a combination of them installed on the system that in operation causes or cause the system to perform the actions. One or more computer programs can be configured to perform particular operations or actions by virtue of including instructions that, when executed by data processing apparatus, cause the apparatus to perform the actions.

The foregoing and other embodiments can each optionally include one or more of the following features, alone or in combination. For instance, in some aspects generating, by the switch proxy controller and based on the second identifier indicated in the first reply, a second reply that indicates the first identifier includes determining that information in the first reply corresponds to a rule programmed by a second switch fabric and generating the second reply to not include the information in the first reply determined to correspond to the rule programmed by the second switch fabric.

In certain aspects, determining that information in the first reply that corresponds to a rule programmed by a second switch fabric includes determining that the information in the first reply that corresponds to the rule programmed by the second switch fabric indicates a rule with a number that is within a range assigned to rules from the second switch fabric. In some aspects, determining that the information in the first reply that corresponds to the rule programmed by